Generalized Suffix Tree based Multiple Sequence Alignment for Service Virtualization

Jean-Guy Schneider and Peter Mandile
Faculty of Science, Engineering and Technology
Swinburne University of Technology
P.O. Box 218, Hawthorn, VIC 3122, AUSTRALIA
{jschneider,pmandile}@swin.edu.au

Steve Versteeg
CA Labs
Level 2, 380 St. Kilda Rd
Melbourne, VIC 3004, AUSTRALIA
steven.versteeg@ca.com

arXiv:1606.01593v1 [cs.SE] 6 Jun 2016

***Abstract*—Assuring quality of contemporary software systems is a very challenging task due to the often large complexity of the deployment environments in which they will operate. Service virtualization is an approach to this challenge where services within the deployment environment are emulated by synthesising service response messages from models or by recording and then replaying service interaction messages with the system. Record-and-replay techniques require an approach where (i) message prototypes can be derived from recorded system interactions (*i.e.* request-response sequences), (ii) a scheme to match incoming request messages against message prototypes, and (iii) the synthesis of response messages based on similarities between incoming messages and the recorded system interactions. Previous approaches in service virtualization have required a multiple sequence alignment (MSA) algorithm as a means of finding common patterns of similarities and differences between messages required by all three steps.**

In this paper, we present a novel MSA algorithm based on Generalized Suffix Trees (GSTs). We evaluated the accuracy and efficiency of the proposed algorithm against six enterprise service message trace datasets, with the proposed algorithm performing up to 50 times faster than standard MSA approaches. Furthermore, the algorithm has applicability to other domains beyond service virtualization.

***Keywords*—service virtualization; service emulation; multiple sequence alignment; enterprise systems; protocol modelling;**

I. Introduction

Organizations today are more reliant than ever on an IT infrastructure to deliver their services, and the software systems providing these services are increasingly interconnected and interdependent. Ensuring quality and correct interoperation of these systems is paramount in achieving uninterrupted business operations in these organizations. However, these software systems are mainly developed in isolation without ready access to testing environments that truly reflect the complexity of real-world deployment environments. As a consequence, we are faced with the increasing risk that unidentified flaws cause a cascade of failures across multiple systems, bringing down an entire IT infrastructure and causing severe interruptions to business operations.

One popular approach that developers use to test their application's dependence on other systems is to install the other systems on virtual machines [1]. However, virtual machines are, in general, time consuming to configure and maintain, and the configuration of the systems running on virtual machines is likely to be different to the real deployment environments. An alternative approach is *service emulation* or *service virtualization* where models of services are emulated – sometimes into the many thousands of service instances – to provide more realistic scale and less complicated configuration [2]. However, service emulation often relies on system experts explicitly modelling the target services and hence require detailed knowledge of the underlying message protocols and structure, respectively. This is often infeasible if the required knowledge is unavailable.

We have been working on an automated approach to service emulation which uses no explicit knowledge of the services, their message protocols and structures, but solely relies on recordings of interactions between a system under test and its environment services [3], [4]. One key aspect of the approach is to identify similarities between messages and exploit these similarities in response generation. Commonly used *multiple sequence alignment* (MSA) techniques from biology [5], [6] were used to do so – refer to [4] for the details of the approach.

However, our experiments have shown that for larger clusters of messages, the memory required to perform a multiple sequence alignment as well as the corresponding computation time increased significantly. Whereas in bioinformatics the main concern is to align few (generally less than 10), but *long* sequences of amino acids, our approach requires the alignment of many (100s to thousands), but *shorter* messages. MSA algorithms optimized for the biological use case may therefore be unsuited for the service emulation domain.

The main contribution of this paper is a novel, memory and time efficient MSA algorithm based on *Generalized Suffix Trees* [7] that addresses these shortcomings. The efficiency and accuracy of the algorithm is evaluated using a set of experiments with typical enterprise system messaging protocols, including LDAP, a binary mainframe protocol, and SOAP services, and compare it with the standard MSA algorithm used in our prior work.

The rest of this paper is organized as follows: we further motivate our work on a novel multiple sequence alignment (MSA) algorithm suitable for the service emulation domain in Section II, followed by a discussion of other MSA algorithms in Section IV. In Section V, we present our novel, GST-based multiple sequence alignment approach. The results of our experiments in evaluating the effectiveness and accuracy of the new approach are presented and discussed in Section VI.

Finally, in Section VII, we summarize the main findings of this paper and give an outlook for future work.

II. Background

The main goal of our work is to produce an emulation environment for enterprise system testing that does not rely on a priori knowledge of the message protocols and structure used by the environment services, but uses message trace recordings collected a priori to produce a response on behalf of a service when invoked by a system under test (SUT) at run-time.

The main idea behind our approach is that if a request sent to an environment service by a SUT is very similar to a recorded request for this service (having a suitable notion of "similarity"), then the response to this request is expected to be similar to the corresponding previously recorded response. Hence, identifying the differences between the incoming and previously recorded requests should give us a good indication how the corresponding recorded response can be altered in order to synthesize a matching response.

In [3], we used an approach where an incoming message was compared with *all* recorded messages in order to determine the most similar one. We used a normalized *edit distance* [8] as the similarity measure. Although this approach produced very accurate responses, it does not scale well to large transaction libraries. Using a dynamic programming approach, the edit distance between two messages m_1 and m_2 with corresponding lengths m_1^l and m_2^l has a time complexity of $\mathcal{O}(m_1^l m_2^l)$. Hence, if we have a transaction library with n recorded messages and an average message length of $\overline{m}$, then finding the most similar recorded request has a time complexity of $\mathcal{O}(n^2\overline{m}^2)$.

In order to improve efficiency, we extended to original approach by clustering the trace recordings into groups of similar messages (ideally of the same *operation type*) and then formulate a single representation for the request messages for each cluster [4]. This accelerated run-time performance by enabling incoming requests from the system under test to be compared only to the cluster representations, rather than the entire transaction library. However, choosing the cluster *centroid*[1] request as the representative resulted in a decreased accuracy of the generated responses as the information from the other requests in the cluster was discarded.

Therefore, we need an approach that addresses both concerns, efficiency and accuracy, by generating cluster *prototypes* which capture the common features of the range of requests in each cluster, but also preserves their variability as much as possible.

In general, messages of typical enterprise-level protocols contain a mix of structural information (which is mostly *identical* for the same operation type) as well as payload information that varies from message to message. Hence, it should be possible to "summarize" the requests within each cluster using a suitably crafted *regular expression* [9] containing constant strings for structural information and patterns for payload information, respectively. As an example, consider

[1]The centroid is the transaction with the minimized total distance from the other transactions in the cluster.

Cluster	Request Message
Search	{id:1,op:S,sn:Smith} {id:275,op:S,sn:Miller} {id:13,op:S,sn:Wilson} {id:2273,op:S,sn:Mandile} {id:490,op:S,sn:Schneider}
Add	{id:24,op:A,sn:Schneider,mobile:123456} {id:4287,op:A,sn:William} {id:3206,op:A,sn:Turner,gn:Samuel,Postcode:34589}

TABLE I: Sample search and add operations

the messages in the search cluster given in Table I. They all conform to the regular expression

`{id:[1-9][0-9]*,op:S,sn:[A-Z][a-z]+}`

assuming that message identifiers do not start with a '0' and can have an arbitrary length and all search strings start with an uppercase character. Similarly, all messages of the add cluster conform to a slightly more complex regular expression. One of the key advantages of regular expressions is that they generally match a broader range of messages than just the ones they were extracted from, as long as these message follow the same basic message structure.

Initial experiments with manually constructing cluster prototypes for LDAP [10] showed that indeed, all clusters prototypes could be expressed as regular expressions. This leads us to the question how to best extract a regular expression from a given set of input sequences.

Existing approaches for the generation of regular expressions from a input samples can be classified into two broad categories: either they generate a finite automaton for each input sequence and then apply a number of rules to merge the individual automata into a single automaton, or they use a multiple sequence alignment (MSA) based approach.

Watson's approach [11], for example, falls into the first category. He presented a semi-incremental algorithm for constructing minimal acyclic deterministic finite automata for a given set of sequences. These automata can then be translated into regular expressions. However, the generated automata only exploit common parts at the beginning and end of the input sequences (in our search cluster example, '`{id:`' and '}'), but not in-between, resulting in rather complex automata. Also, the resulting automata only accept sequences from the given input, but do not accept any sequences not covered by the input, which is of rather limited use for our purpose.

MSA-based approaches, such as for example the one presented by Tang *et al.* [12], perform an alignment of all sequences, identify the overlapping (or aligned) sub-sequences as constant strings of the resulting regular expression, and extract suitable patterns for the gaps in-between. Consider the multi-sequence alignment of the messages in the search cluster given in Table I (with ⋆ denoting the *gap symbol* introduced during the alignment):

```
{id:1⋆⋆⋆,op:S,sn:Smith⋆⋆⋆⋆}
{id:275⋆,op:S,sn:Miller⋆⋆⋆}
{id:13⋆⋆,op:S,sn:Wilson⋆⋆⋆}
{id:2273,op:S,sn:Mandile⋆⋆}
{id:490⋆,op:S,sn:Schneider}
```

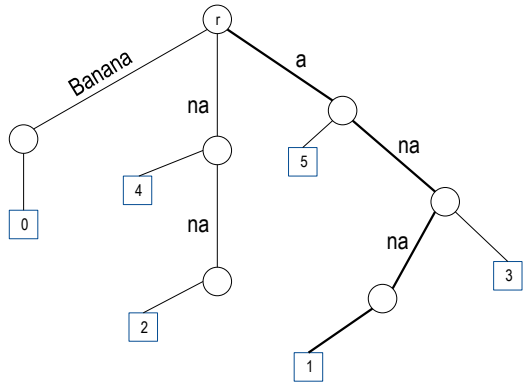

Fig. 1: Suffix Tree for the word *Banana*.

The alignment results in three overlapping sub-sequences (*i.e.* '{id:', ',op:S,sn:' and '}'), the same constant sub-sequences that were used to form the regular expression shown above.

For the rest of this work, we solely focus on identifying the overlapping (or aligned) sub-sequences of a set of messages in order to identify the constant parts of regular expression-based cluster prototypes. Finding suitable patterns for the gaps between the constant strings that appropriately reflect the nature of the payload is a topic of an ongoing investigation.

III. Suffix Trees

In this section, we will briefly introduce a number of concepts required to illustrate the proposed new multi-sequence alignment algorithm, especially the notion of a *Generalized Suffix Tree* (GST) our approach is heavily based upon.

A. Sequences and Suffixes

We start with the notion of the most basic building block for our study, the set of *characters*, denoted by $\mathcal{C}$. We require equality and inequality to be defined for the elements of $\mathcal{C}$. For the purpose of our study, $\mathcal{C}$ will most likely comprise of the set of valid Bytes that can be transmitted over a network or the set of printable Characters as a dedicated subset. A word (or *sequence*) s is a non-empty, finite sequence of characters $c_0c_1c_2\ldots c_{n-1}$ with $c_i \in \mathcal{C}$, $0 \leq i < n$. The *length* of a sequence s, denoted by s^l, is equal to the number of characters, that is, $s^l = n$. Throughout the rest of this section we will use the terms word and sequence interchangeably.

Given a word $s = c_0c_1c_2\ldots c_{n-1}$, we define a *sub-word* as a non-empty sub-sequence of characters from s starting at index i and having length l as $sw_s(i,l) = c_ic_{i+1}c_{i+2}\ldots c_{i+l}$. We require $l > 0$ and $i+l \leq n$. A special kind of subword is a *suffix*, a sub-sequence $c_ic_{i+1}c_{i+2}\ldots c_{n-1}$ that ends with the last character c_{n-1} of the sequence s.

B. Suffix Tree

A suffix tree for a sequence s, denoted by $ST(s)$, is a tree that contains all *suffixes* of s, that is, the sub-sequences c_{n-1}, $c_{n-2}c_{n-1}$ etc. [7]. More specifically, a suffix tree is an *edge-labelled tree* containing (i) n labeled leaf nodes and (ii) up to $n+1$ non-leaf (or *branching*) nodes. The concatenation of the edge labels of the path from the root node of the tree to the leaf node with label i defines the suffix $c_ic_{i+1}\ldots c_{n-1}$.

In order to minimize the number of branching nodes in the tree, a branching node b contains *at most* one outgoing edge labelled $c\alpha$ for a given character $c \in \mathcal{C}$ and some (possibly empty) sequence α with both, α and $c\alpha$ being sub-words of the sequence s. The edges to leaf nodes do not contain any labels.[2] Similarly, the incoming edge to a branching node $b \neq r$ must contain a non-empty label β such that β is a sub-word of s and for each outgoing edge of b to a branching node b' (with edge label γ), $\beta\gamma$ is also a sub-word of s. In order to guarantee that a suffix tree for a sequence with n characters has no more than $n+1$ branching nodes, each branching node must have either an outgoing edge to at least one leaf node or to two branching nodes, respectively.

Consider the suffix tree for *Banana* as given in Figure 1. Branching nodes are depicted as circles, leaf nodes as squares. The suffix tree has a uniquely identifiable root node, denoted by r, 6 further branching nodes, and 6 leaf nodes (labelled 0 to 5). The path highlighted in Figure 1 (from r to leaf node with label 1) denotes the suffix starting at index 1, that is, '*anana*'.

It is well known that suffix trees can be constructed in linear time and with linear space [7], [13].

C. Generalized Suffix Tree

A *generalized suffix tree* (or short GST) generalizes the concept of a suffix tree to more than one sequence, that is, it contains all suffixes for sequences $s_0, s_1, \ldots, s_{m-1}$ (with $m > 1$). In order to distinguish which sequence a particular suffix belongs to, leaf nodes contain two labels: one for the sequence, and one for the start index of the suffix in the corresponding sequence. As a consequence, a GST for the sequences $s_0, s_1, \ldots, s_{m-1}$ contain $s_0^l + s_1^l + \ldots + s_{m-1}^l$ leaf nodes. The other properties as given for a suffix tree still hold. A GST can be constructed in $\mathcal{O}(s_0^l + s_1^l + \ldots s_{m-1}^l)$ time and space, respectively [13].

We assign a *colour set* to each branching node in a GST [15]. Each sequence has a unique colour, and we assign this colour to all leaf nodes that represent a suffix for this sequence. The colour set of each branching node is the union of the colour set(s) of all of its children. As a consequence, the colour set of the root node r contains the colours of all sequences. The colour set of a branching node can be determined in a top-down tree construction process and does not need to be re-computed bottom-up once a GST is fully constructed. Hence, the computation of the colour sets does not add to the algorithmic complexity of the tree construction.

The reader may note that from a sequence alignment perspective, of special interest are branching nodes that have the same colour set as the root node [15]. We denote such branching nodes as *fully coloured*. The concatenation of the edge labels of the paths from r to these nodes correspond to sub-sequences that are common to *all* sequences of the GST.

[2] Some references (*e.g.* [13], [14]) use a special character \$ to label edges to leaf nodes, indicating a dedicated *end-of-word* character. For the purpose of this presentation, we have omitted edge labels to leaf nodes.

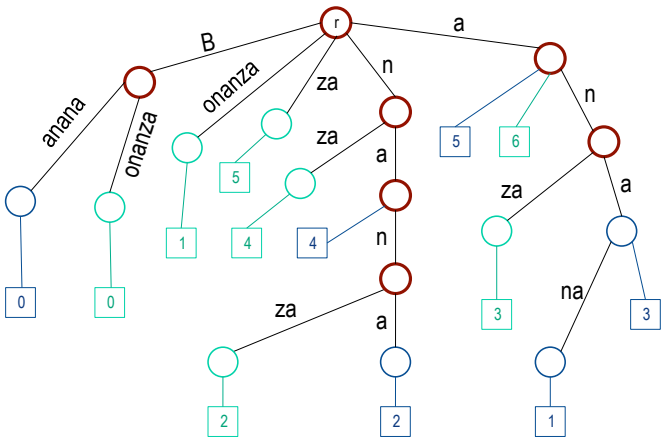

Fig. 2: Generalized Suffix Tree for *Banana* and *Bonanza*.

As an example, consider the generalized suffix tree for the two sequences *Banana* and *Bonanza* as given in Figure 2. In order to enhance readability, the leaf nodes for *Banana* and *Bonanza* are denoted by blue and green squares, respectively. All fully coloured nodes are highlighted in red. As we can infer from Figure 2, the two sequences contain the common sub-sequences `B`, `n`, `na`, `nan`, `a`, and `an`, respectively.

D. Multi Sub-words

The concatenation of the edge labels to all fully coloured branching nodes give us the *values* of the common sub-sequences, but not their respective *positions* (or start indices) in the sequences of a GST. The start indices are defined by the labels of all leaf nodes in the corresponding sub-trees of each of the fully coloured branching nodes. For example, the bottom-most fully coloured branching node in Figure 2 represents the common sub-sequence `nan`. Its subtree contains leafnodes for index 2 for both sequences, hence `nan` starts at index 2 in *Banana* and *Bonanza*, respectively.

If there is more than one leaf node for a specific colour (or sequence), then the corresponding sub-sequence appears *multiple times* in this sequence, and multiple combinations can be found across all sequences. For example, the sub-sequence `a` appears three times in *Banana* and twice in *Bonanza*, hence a total of 6 combinations for `a` can be found. More generally, if we have a common sub-sequence α that appears m times each in n sequences, there is a total of m^n possible combinations. Therefore, MSA approaches (such as the one proposed in [14]) that rely on generating all possible combinations during an alignment process, are unlikely to scale to a large number of sequences.

In order to avoid enumerating all possible combinations for a given common sub-sequence, we introduce the notion of a *multi sub-word* (or MSW). A multi sub-word is a combination of the value of a common sub-sequence, together with a set of occurrences thereof across all sequences. We use the notation

$$[\ \mathtt{a} - \{\ 0@1\ 0@3\ 0@5\ 1@3\ 1@6\ \}\ ]$$

for multi sub-words, in this case, to indicate that sub-sequence `a` appears at indices 1, 3 and 5 in sequence 0 (*i.e. Banana*) and at indices 3 and 6 in sequence 1 (*i.e. Bonanza*), respectively.

IV. Multiple Sequence Alignment

Multiple sequence alignment (MSA) was first used to align three or more biological sequences to reveal their structural commonalities [16]. In contrast to pair-wise sequence alignment algorithms such as Needleman-Wunsch [17] or Smith-Waterman [18] that have a quadratic time and space complexity, creating an optimal alignment for multiple sequences is an NP-complete problem [19]. Therefore, a number of heuristic techniques were proposed, with *ClustalW* [6] probably being the most prominent one.

A brief overview of the *ClustalW* algorithm is as follows:

1) All $n(n-1)/2$ possible pairs of sequences are aligned by using a standard Needleman-Wunsch algorithm [17] in order to calculate their *similarity*, resulting in an $n \times n$ similarity matrix.
2) A *guide tree* is constructed from the similarity matrix by applying a neighbour-joining clustering algorithm [20].
3) The tree is then used to guide a progressive pair-wise alignment process by traversing the tree from the leave nodes to the root.

ClustalW has a polynomic algorithmic complexity, but the first step alone is $\mathcal{O}(n^2m^2)$ with m being the average sequence length.

In order to better exploit similarities between sequences, a number of GST-based MSA algorithms have been proposed (*cf.* [5], [14], [21], [22] just to name a few). Höhl *et al.* [14], for example, proposed a MSA algorithm that (i) creates a GST for the sequences to be aligned, (ii) identifies all nodes in the GST that adhere to a property equivalent to the notion of fully coloured, and then (iii) generates all possible combinations of sub-sequences that are *maximal*. A sub-sequence is maximal if it cannot be extended either to the left or right without having at least one character that is not common across all sequences. The collection of maximal common sub-sequences forms the basis of a directed, edge-weighted graph (V, E) with the maximal common sub-sequences as vertices. A weighted, directed edge is added between s_1 and s_2 if and only if s_2 is "to the right" of s_1 in all sequences. The weight of this edge is the length of s_1. The best possible alignment for all sequences is then defined by a chain through the graph with maximal weight-length [14].

Consider the application of this algorithm to the two sequences `ADCxzDCxBAx` and `DCxAzDCxpxBA`. This leads to the following, non-optimal alignment:

```
ADCx⋆zDCxBAx⋆⋆
⋆DCxAzDCxp⋆xBA
```

There is a total of 24 different common sub-sequences, only 12 thereof are both right- and left-maximal. However, as `zDCx` and `xBA` overlap in the first sequence (they do not overlap in the second sequence), `xBA` is *not* to the right of `zDCx`, only the last `x` in both sequences are. Hence, as `BA` is not a maximal sub-sequence and the algorithm does not allow for trimming of partially overlapping sub-sequences, `BA` is not aligned, only `x` is. Therefore, besides the problem of a possible very large number of all common sub-sequence combinations (*cf.* Section III-D), this algorithm does not cater well for partial overlaps and may generate sub-optimal alignments.

[1] Create a generalized suffix tree for the sequences to be aligned.
[2] Identify all *fully-coloured branching nodes* and create a corresponding collection of multi sub-words.
[3] Choose the "best" multi sub-word MSW from this collection.
[4] Choose the "best" common sub-sequence CSW from the multi sub-word MSW. CSW will act as *anchor* for the first alignment.
[5] For each multi sub-word in the collection, remove any sub-sequences that *fully overlap* with CSW.
[6] For each multi sub-word in the collection, trim any sub-sequences that *partially overlap* with CSW. Due to the trimming, this step can either add new multi-subwords or add sub-sequence information to existing multi sub-words.
[7] Split the collection of multi sub-words into two collections: one that is *entirely to the left* and one that is *entirely to the right* of the anchor CSW.
[8] Recursively apply steps 3 and 7 to both, the left and right collections. Terminate the recursion at Step 3 if the collection of multi sub-words is empty.

Fig. 3: The main steps of the *Mandile-Schneider* algorithm.

Common to all MSA algorithms we investigated is that they were designed to align *few*, but possibly long sequences. This contrasts with our problem of creating prototypes for clusters that contain *many*, but possibly short(er) sequences. This mismatch motivated the new MSA algorithm presented in the following section.

V. MANDILE-SCHNEIDER ALGORITHM

In order to address the scalability issues of existing MSA algorithms, we defined a new approach that (i) exploits the benefits of using a GST to identify common sub-sequences and (ii) uses *divide-and-conquer* to align segments that progressively become smaller. In line with other alignment algorithms (such as Needleman-Wunsch [17] and Smith-Waterman [18]), we call this algorithm the *Mandile-Schneider* algorithm. The guiding principles behind the algorithm are that we want to (i) exploit commonalities between sequence as early as possible, (ii) avoid the enumeration of all possible combinations of a common sub-sequence (either entirely or as much as possible) as this can become very computationally intensive, and (iii) give precedence to long common sub-sequences over shorter ones when applying the divide-and-conquer principle. Most importantly, though, the approach should be fit for purpose for our specific needs, that is, the alignment of many, but short sequences. The main steps of the new algorithm are outlined in Figure 3.

The result of the algorithm is a list of common sub-sequences that defines all anchors of the alignment. A gap between two consecutive sub-sequences is an area where no alignment was found.

The reader may note that the identification of the "best" multi sub-word and common sub-sequence in steps 3 and 4, respectively, is left open. These are two *variation points* in the algorithm where different strategies can be adapted to identify the "best" anchor. In line with the guiding principles mentioned earlier, we use a *biggest-left-most-common-subword* strategy for the majority of this work, that is, the *left-most* common sub-sequence in the *largest* multi sub-word is chosen as the anchor for each alignment step.[3]

Let's illustrate the main steps of the *Mandile-Schneider* algorithm on the two sequences `ADCxzDCxBAx` and `DCxAzDCxpxBA` used in the discussion of the graph search-based approach by Höhl *et al.* [14]:

- **Step 1 and 2:** the GST of the two sequences has 7 fully coloured branching nodes, leading to the following collection of multi sub-words:
  `[ A – { 0@0 0@9 1@3 1@11 } ]`
  `[ BA – { 0@8 0@10 } ]`
  `[ Cx – { 0@2 0@6 1@1 1@6 } ]`
  `[ DCx – { 0@1 0@5 1@0 1@5 } ]`
  `[ x – { 0@3 0@7 0@10 1@2 1@7 1@9 } ]`
  `[ xBA – { 0@7 1@9 } ]`
  `[ zDCx – { 0@4 1@4 } ]`
- **Step 3 and 4:** the longest multi sub-word is `zDCx`, the left-most common sub-sequence thereof starts at index 4 in both words. This becomes the first anchor.
- **Step 5:** the removal of any overlaps with `zDCx` at index 4 in both words reduces the collection of multi sub-words to the following:
  `[ A – { 0@0 0@9 1@3 1@11 } ]`
  `[ BA – { 0@8 0@10 } ]`
  `[ Cx – { 0@2 1@1 } ]`
  `[ DCx – { 0@1 1@0 } ]`
  `[ x – { 0@3 0@10 1@2 1@9 } ]`
  `[ xBA – { 0@7 1@9 } ]`

  `zDCx` overlaps with itself, hence is fully removed, and one occurrence each of `DCx` and `x` is removed.
- **Step 6:** the only partially overlapping multi-subword is `xBA` – it is trimmed back to `BA`. Since we already have a multi sub-word for `BA` at indices 8 and 10, respectively, the trimmed multi sub-word does not further contribute to the collection of multi sub-words.
- **Step 7:** splitting the multi sub-words into what is entirely to the left and right of `zDCx`, respectively, results in the following two collections of multi sub-words:

  Left:
  `[ A – { 0@0 1@3 } ]`
  `[ Cx – { 0@2 1@1 } ]`
  `[ x – { 0@3 1@2 } ]`
  `[ DCx – { 0@1 1@0 } ]`

  Right:
  `[ A – { 0@9 1@11 } ]`
  `[ BA – { 0@8 0@10 } ]`
  `[ x – { 0@10 1@9 } ]`

[3]The current implementation chooses the *first* multi sub-word if more than one largest is found.

Please note that A and x occur in both the left and right collections as there is one occurrence each on either side of the anchor zDCx.

- **Recursion on *left*:** DCx is the longest multi sub-word in *left*, the left-most common sub-sequence thereof starts at indices 1 and 0, respectively. The removal of overlaps only leaves A, but since A *crosses over* the anchor DCx (the occurrence of A in the first sequence is left of the anchor, but to the right in the second sequence), the resulting left and right collections in Step 7 are empty, and the recursion for both terminates in Step 3.
- **Recursion on *right*:** BA is the longest multi sub-word in *right*, the left-most common sub-sequence thereof starts at indices 8 and 10, respectively. Similar to the recursion on *left*, the removal of overlaps leaves one multi sub-word only (*i.e.* x) that crosses over the anchor, hence the resulting left and right collections in Step 7 are empty, and the recursion terminates as well.
- **Result:** we get three anchors DCx, zDCx, and BA that define the following, optimal alignment:

ADCx⋆zDCx⋆⋆BAx
⋆DCxAzDCxpxBA⋆

We implemented the *Mandile-Schneider* algorithm as a set of Java classes and also integrated a number of packages from BioJava[4] needed for the evaluation presented in the next section. During the implementation, we applied a number of optimizations in order to improve both, the run-time performance as well as the memory footprint. However, there is room for further optimizations. A packaged version of the software is available from http://quoll.ict.swin.edu.au/doc/ms/.

VI. Evaluation

In this section, we present the experiments that we conducted to evaluate the accuracy and effectiveness of the novel MSA approach presented in the previous section and discuss the results of our experiments. More specifically, we introduce our experimental setup as well as our evaluation criteria in Section VI-A. In Section VI-B, we present and discuss the results of our evaluation, especially from a resource usage and accuracy perspective. We discuss limitations of our current approach and identify possible areas of future improvements in Section VI-C. Finally, we present an initial evaluation of the new MSA algorithm on biological data in Section VI-D.

A. Experimental Setup

Although one of the main aims of our work is to generate regular expressions to summarize message clusters for unknown or ill-specified protocols, for evaluation purposes, we used three protocols where the precise message structures for the various operation types are known: the Lightweight Directory Access Protocol (LDAP) [10], the Simple Object Access Protocol (SOAP) [23], and IBM Information Management System (IMS) [24] which is widely used on mainframe computers.

The recorded LDAP messages used for this evaluation consisted of the core LDAP operations ADD, SEARCH, and MODIFY applied to CA Technologies' *DemoCorp* sample directory [25]. We used a decoder to translate the corresponding messages into a text format, but kept a subset of the SEARCH operations in their binary format.

Protocol	Operation	Type	No. Msgs.	Median Length
LDAP	Add	text	851	378
LDAP	Modify	text	334	239
LDAP	Search	text	621	313
SOAP	Withdraw	text	160	243
LDAP	Search	binary	606	73
IMS	Update	binary	200	231

TABLE II: Protocols and operations used for evaluation.

The SOAP messages used for evaluation purposes were generated based on a recording of a banking example using the LISA tool [26]. The protocol consists of 7 different request types, each with a varying number of parameters, encoding "typical" transactions one would expect from a banking service. We chose to use the messages from the WITHDRAWMONEY operation for this evaluation.

The IMS messages used were recorded interactions between a client tool and an IMS phone book demonstration server. IMS messages consist of some header data followed by the record structure encoded as *fixed width fields*. The demonstration server supported four different transaction types including: ADD, UPDATE, DISPLAY and DELETE. In addition the recording included IMS acknowledgement messages. For the purpose of the evaluation, we used the messages of the UPDATE transaction type.

A summary of the protocols, operations, number of messages as well as the median message length (measured in number of characters) is given in Table II. The raw input data used for the evaluation is available at http://quoll.ict.swin.edu.au/doc/ms/.

For the purpose of the evaluation, we were interested in the (i) resource usage of the new approach and (ii) the quality and accuracy of the generated alignments. In order to compare the obtained results, we used the ClustalW [6] implementation of BioJava 4.0[4] as a comparison benchmark.

Resource usage was measured in both *total computation time* for an alignment (excluding the time to read the input and analyzing the generated alignment) as well as the *memory usage* as reported by the Java Virtual Machine.[5] The *quality* of each generated alignment is measured using an *edit distance* [8] as well as the total number of overlapping (*i.e.* aligned) characters. With regards to accuracy, for each of the operation types of the six protocols/operations used, we identified a number of sub-sequences (representing specific structural information of the corresponding operation types) and checked whether these sub-sequences were aligned across all messages.

For each of the six input sets used for evaluation, the input files with the corresponding messages were *randomly* ordered in order to avoid any bias from the corresponding recordings. Therefore, in order to create an alignment for n messages, we simply used the n first lines in the corresponding file as input.

[4] http://www.biojava.org/

[5] This was measured by subtracting the free memory from the total memory directly after the generation of the alignment.

For each protocol/operation, we generated alignments for 2, 5, 10, 15, 20, 25, 50, 100, and 200 messages as well as the maximal available number of messages in order to illustrate the scalability of both, the proposed approach as well as the comparison benchmark. Each alignment computation was repeated *25 times* in order to mitigate against any "random" effects from other processes running on the computing platform.

All experiments were run on an iMac with a 2.7 GHz Intel Core 5 CPU, 8GM of RAM, running Mac OS X 10.8.5 and Java 1.8.

B. Evaluation Results

The results of our evaluation experiments are summarized in Table III: for each combination of operation type/number of messages to align it lists: the median total time taken to generate an alignment (in milli-seconds), the median memory usage (in Megabytes), the edit distance of the resulting alignment, and the number of overlapping characters (*i.e.* the sum of the lengths of all common sub-sequences in the alignment). For the proposed *Mandile-Schneider* approach, we also list the number of multi sub-words that were extracted from the generalized suffix tree (not applicable to ClustalW). The final table column, titled 'Speed-Up', shows the performance increase of *Mandile-Schneider* compared to ClustalW.

1) Resource Usage: Table III illustrates the differences in time complexity for the proposed algorithm compared to ClustalW, respectively: for all data points, the total time taken for *Mandile-Schneider* is significantly better than the one for ClustalW. The speed-ups vary from 3.78 for 2 sequences in LDAP Search to 55.63 for 851 sequences in LDAP Add.

A brief statistical analysis demonstrated a strong quadratic correlation ($R^2 > 0.99$) between the number of sequences and the total alignment time for ClustalW for all of the six protocols/operations. This is not surprising as for n sequences, ClustalW needs to compute the similarities of all $n*(n–1)$ possible pairs in order to build the guide tree. The corresponding (quadratic) coefficients varied from 0.097 for the binary LDAP Search to 2.50 for LDAP Add. However, even for binary LDAP Search (where the individual messages are short compared to the other five case examples), the quadratic correlation was much stronger than a corresponding linear relationship. On the other hand, our proposed algorithm exhibits a linear correlation for the number of sequences/computation time relationship for LDAP Modify, SOAP, binary LDAP Search and IMS, respectively, whereas the other two protocols/operations showed a quadratic correlation (but with coefficients around 0.03 for the quadratic term).

We can conclude that the proposed algorithm is faster than the ClustalW implementation of BioJava 4.0, and the more sequences need to be aligned, the greater the speedup. To some extent this is not surprising given the algorithmic complexity of ClustalW. However, we will need to run further experiments on different kind of protocols and/or operations and perform both, a detailed statistical analysis as well as an algorithmic complexity analysis to get further evidence of a possible linear time complexity of the proposed algorithm.

The results for the memory usage are unfortunately not as consistent. Consider the results for all sequence combinations for LDAP modify as given in Figure 4.

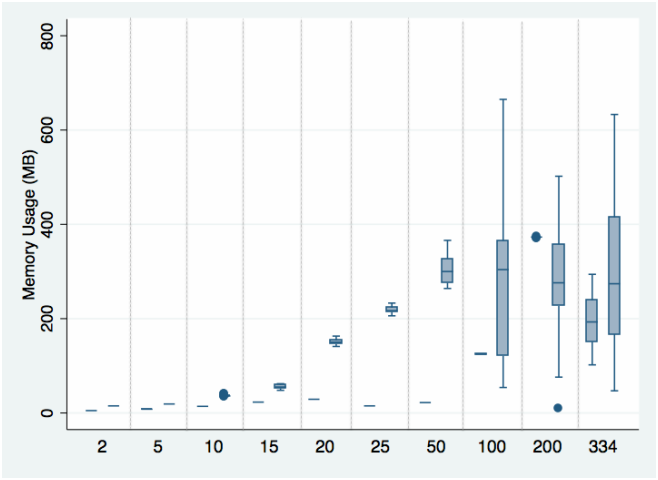

Fig. 4: Memory usage for LDAP Modify: *Mandile-Schneider* on the left, ClustalW on the right. Number of sequences to be aligned on the x-axis.

For ClustalW, once the similarity between two sequences is established, the matrix required to compute the similarity is not needed any more and hence can be garbage collected. Therefore, the larger the number of sequences, the more likely the garbage collector will kick in, but as Figure 4 illustrates, this does not happen in a consistent manner (*e.g.* demonstrated by the large spread for 100 sequences). Similarly, for the proposed new approach, once all multi sub-words are identified, the generated GST is not needed any more and can be garbage collected. We suspect this has happens as the median memory usage for all 334 sequences for LDAP Modify is lower than for 200 sequences. However, without a more accurate measurement of the maximal memory usage we cannot really draw too many conclusions from the experiments.

2) Accuracy: To the best of our knowledge, the generated alignments of both approaches do separate message structure information from payload, precisely as was required. This was verified by a combination of manual inspection and hand-crafted regular expressions that were run over the generated alignments. With the exception of IMS Update (which will be discussed in more detail in the next section), the edit distances of *Mandile-Schneider* are no more than 10% worse than ClustalW. The differences in edit distance between the two alignment approaches is mainly due to how the payload is aligned – structure information is aligned consistently. Whereas ClustalW minimizes the number of gaps, *Mandile-Schneider* maximizes the total number of overlapping characters which may result in larger gaps between aligned subsequences, and, consequently, a larger edit distance. Initial results using the generated alignments to create regular expressions as cluster prototypes (the main aim of our investigation) indicate that the alignments are accurate enough to do so, but a further, more detailed analysis is required to substantiate this claim. This is the topic of an ongoing investigation.

An interesting observation of our experiments is that the total number of *multi sub-words* generated by the corresponding generalized suffix tree is quite low. In case of LDAP Add, for example, even the GST for all 851 messages does not define

Protocol	#Seq.	Mandile-Schneider Time(ms)	Mem.(MB)	Edit-Dist.	#Chars	#MSWs	ClustalW Time(ms)	Mem.(MB)	Edit-Dist.	#Chars	Speed-Up
LDAP Add	2	82	9	119	290	376	309	23	107	288	3.77
	5	109	15	579	267	368	719	14	481	256	6.60
	10	158	27	1279	263	389	1324	40	1238	257	8.38
	15	212	20	2089	263	391	2041	121	2053	253	9.63
	20	266	11	3019	263	397	2959	69	2870	254	11.12
	25	310	42	3945	263	401	4198	304	3662	254	13.54
	50	523	85	8635	263	407	10656	218	8402	252	20.37
	100	1069	190	18475	263	416	33103	260	17483	252	30.97
	200	2674	206	37511	263	430	113122	192	36738	250	42.30
	500	12925	416	97875	263	441	667836	260	93450	246	51.67
	851	33701	761	166967	263	447	1874737	316	159297	246	55.63
LDAP Modify	2	60	5	158	161	249	222	15	147	164	3.70
	5	72	8	849	150	230	448	19	614	147	6.22
	10	118	14	2109	133	174	792	36	1497	126	6.71
	15	145	23	2602	132	178	1146	55	2448	121	7.90
	20	169	29	3579	132	178	1545	150	3053	110	9.14
	25	207	15	4546	132	180	2222	217	3858	110	10.73
	50	349	22	9683	132	189	5376	300	8394	106	15.40
	100	608	125	20292	132	199	14946	304	18566	109	24.58
	200	1584	373	43660	132	208	50269	276	41044	109	31.74
	334	2786	193	73098	132	214	131497	274	64415	112	47.20
LDAP Search	2	69	6	36	284	392	261	18	29	282	3.78
	5	95	11	157	283	404	502	11	121	282	5.28
	10	140	23	649	274	394	825	34	563	272	5.89
	15	187	12	1011	272	406	1304	59	881	268	6.97
	20	218	30	1361	272	411	1819	173	1185	268	8.34
	25	266	24	1893	272	423	2560	44	1678	268	9.62
	50	468	42	3808	272	434	6447	115	3384	269	13.78
	100	879	44	7792	272	449	19952	280	6893	268	22.70
	200	2114	181	17357	272	465	71828	322	15602	268	33.98
	500	9819	490	45473	272	487	414885	341	41007	268	42.25
	621	14018	536	56537	272	490	636085	354	50982	268	45.38
SOAP Withdraw	2	57	5	11	234	349	224	14	10	233	3.93
	5	71	7	40	230	336	398	18	40	230	5.61
	10	97	11	93	230	334	591	61	93	230	6.09
	15	129	16	145	230	334	922	94	145	230	7.15
	20	136	22	228	230	334	1240	67	228	230	9.12
	25	158	29	286	230	334	1663	95	286	230	10.53
	50	265	21	685	230	334	4373	568	598	229	16.50
	100	469	87	1410	230	334	13380	295	1226	229	28.53
	160	732	115	2411	230	335	31449	229	2116	229	42.96
LDAP Search (binary)	2	35	3	36	36	42	157	8	35	35	4.49
	5	45	4	137	34	46	197	11	137	34	4.38
	10	54	5	563	33	44	260	24	566	32	4.81
	15	59	5	886	32	46	315	9	888	32	5.34
	20	67	7	1194	32	48	373	31	1194	32	5.57
	25	74	8	1664	32	51	471	29	1667	32	6.36
	50	108	14	3377	32	55	906	83	3343	33	8.39
	100	166	31	6915	32	58	2099	156	6913	32	12.64
	200	291	33	15503	32	59	5785	245	15491	32	19.88
	500	769	97	40316	32	60	27898	578	40353	32	36.28
	606	973	169	48887	32	61	40658	377	48924	32	41.79
IMS Update (binary)	2	69	5	37	210	306	221	13	26	207	3.20
	5	92	9	180	207	299	380	15	90	205	4.13
	10	130	15	450	205	303	549	47	186	204	4.22
	15	168	23	716	205	303	777	62	279	201	4.63
	20	202	31	836	204	303	1098	131	401	202	5.44
	25	243	16	1047	204	303	1352	256	498	203	5.56
	50	386	21	2201	203	304	3332	226	1100	201	8.63
	100	701	100	3976	201	306	9227	256	2105	200	13.16
	200	1718	170	8352	200	301	31009	218	4764	198	18.05

TABLE III: Results of alignment experiment using six different operation types.

more than 447 multi sub-words (*cf.* Table III). On the other hand, the total number of *common sub-sequences* contained in these multi sub-words grow exponentially. For the first two messages, there are 26,339 common sub-subsequences, growing to 13,906,721,076 (*i.e.* more than 13 Billion) for five messages. For 10 messages, the total number of common sub-sequences cannot be represented using a 64 Bit integer any more (*i.e.* it is greater than 10^{18}), demonstrating that the graph-based approach introduced by Höhl *et al.* [14] does indeed not scale beyond small examples as postulated in Section IV.

C. Discussion and Limitations

As illustrated in Table III, there is one protocol/operation type that performs much worse with regards to the edit distance for the proposed algorithm compared to ClustalW: IMS Update. In the following, we briefly discuss one of the main reasons for this significantly worse edit distance.

Consider the following extract of the payload information (surname, given name, area code, phone number) for four messages that each contain a varying number of white spaces

(shown as '␣') in order to ensure a fixed-length message encoding:

```
Garufi␣␣␣␣Meaghan␣␣␣598␣␣␣␣␣␣␣37110
Hoogland␣␣Tamar␣␣␣␣␣575␣␣␣␣␣␣␣14895
Lindall␣␣␣Fabian␣␣␣␣371␣␣␣␣␣␣␣80126
Fern␣␣␣␣␣␣␣Natalie␣␣␣937␣␣␣␣␣␣␣82901
```

The *Mandile-Schneider* algorithm generates the following, non-optimal alignment with an edit distance of 109:

```
Garufi**********␣␣␣␣Meaghan␣␣␣598**␣␣␣␣␣␣␣37**110**
Hoogland␣␣Tamar*␣␣␣␣␣575**********␣␣␣␣␣␣␣␣␣****14895
Lindall␣␣␣Fabian␣␣␣␣371***********␣␣␣␣␣␣␣80**126**
Fern***********␣␣␣␣␣␣␣Natalie␣␣␣937␣␣␣␣␣␣␣82901****
```

The longest common sub-sequence is the seven white spaces between the area code and phone numbers and this sequence is (correctly) aligned first. The longest sub-sequence to its left are four white spaces: between `Garufi` and `Meaghan` for the first sequence, the first four white spaces between `Tamar` and `575` for the second, between `Fabian` and `371` for the third, and finally the first four white spaces between `Fern` and `Natalie`, respectively. Once this sub-sequence is aligned, the rest "in the middle" does not contain any common characters any more and cannot be further aligned.

Please also note that all phone numbers contain at least one number '`1`' and the left-most occurrences thereof are aligned. One may argue that such short common sequences should not be aligned and it would be straightforward to define a threshold to ensure that any alignments shorter than this threshold are ignored. However, the binary LDAP search has a single character '`c`' that indicates the operation type and if we introduce such a threshold then the `c`'s are unlikely to be aligned. Therefore, we decided against introducing such a threshold.

In contrast, ClustalW generates a much better alignment (edit distance of 60), but even this alignment is not optimal:

```
Garufi␣␣␣␣Meaghan␣␣␣598*␣␣␣␣␣␣␣37110
Hoogland␣␣Tamar␣␣␣␣␣575*␣␣␣␣␣␣␣14895
Lindall␣␣␣Fabian*␣␣␣␣371␣␣␣␣␣␣␣80126
Fern␣␣␣␣␣␣␣Natalie␣␣␣937*␣␣␣␣␣␣␣82901
```

Aligning the `1`'s in the phone numbers would introduce new gaps, hence this is omitted, but it is not clear why a gap between `Fabian` and the following white spaces is introduced, but not for `Tamar`.

D. Biological Data

Aligning two (or more) sequences of nucleotide or amino acid residues to identify regions of similarity that may be a consequence of functional, structural, or evolutionary relationships is standard practice in bioinformatics. Hence, in order to test the accuracy of the new alignment algorithm with different input data, we chose to align the first 10,000 amino acids of both, the human and the mouse genome.[6] As a benchmark, we used the alignment generated by the "standard" Needleman-Wunsch algorithm [17].

Figure 5 summarizes the results of this experiment. We only plot an alignment if an aligned sub-sequence is more than

[6]Refer to [22] for specific details about the two genomes.

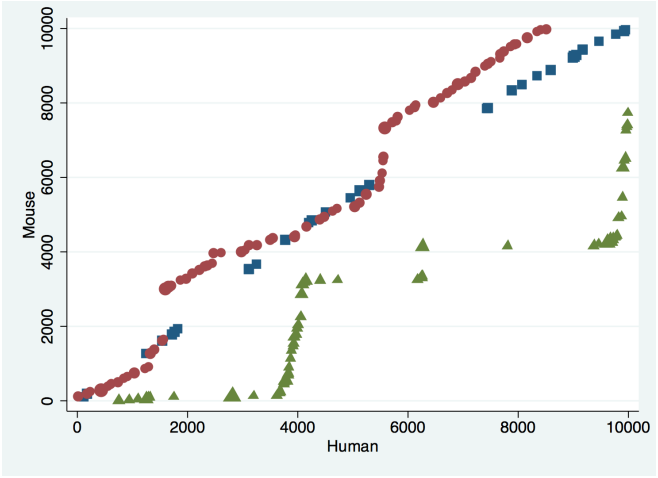

Fig. 5: Alignments of Human and Mouse genomes. Blue squares: Needleman-Wunsch. Green triangles: *Mandile-Schneider* (biggest-left-most-common-subword). Red circles: *Mandile-Schneider* (9 largest multi sub-words).

5 characters (*i.e.* amino acids) long. The size of each symbol relates to the total number of characters in the corresponding sub-sequence. The alignment of the benchmark Needleman-Wunsch algorithm (represented as blue squares in Figure 5) generated only one sub-sequence of 10 characters with 33 sub-sequences (out of a total of 3,479) having a length greater than 5. This indicates that the human and the mouse genomes do not have a large overlap in the first 10,000 amino acids.

The *biggest-left-most-common-subword* strategy does not generate an alignment that is comparable to the one produced by the benchmark Needleman-Wunsch – it is represented by green triangles in Figure 5. The largest alignment is 20 characters long and there are a total of 62 sub-sequences (out of 244) with a size greater than 5. However, the largest alignment is found at index 2,819 in the human genome, but 98 in the mouse, a significant difference between the two. An algorithm that favours size (of sub-sequences) over "location" is probably not a good fit to align the first 10,000 amino acids of the genomes of human and mouse.

Therefore, we devised another strategy to identify the anchor at each step of the alignment. Instead of getting the left-most common sub-sequence of the largest multi sub-word, we identify the n largest multi sub-words, generate all common sub-sequences thereof, and choose the common sub-sequence with the smallest variance in the relative starting indices (normalized by size) as the new anchor. Figure 5 includes the alignment generated by this strategy (represented by red circles) with $n = 9$. The largest alignment is 15 characters long and there are a total of 81 common sub-sequences (out of 1,645) with a size greater than 5.

Clearly, the revised strategy seems to be a much better fit than the original *biggest-left-most-common-subword* strategy. It also shows the flexibility of the proposed approach as new anchor identification strategies can be defined that better fit the specific nature of a given alignment problem at hand.

VII. Conclusions and Future Work

In this paper, we have presented a novel, time and memory efficient Multiple Sequence Alignment (MSA) algorithm based on Generalized Suffix Trees (GSTs). We evaluated the accuracy and efficiency of the new algorithm against six enterprise service message trace datasets, with the proposed algorithm performing up to 50 times faster for large number of sequences than standard MSA approaches. We also evaluated the new approach using a data set from bioinformatics, demonstrating the adaptability of the approach to different domains.

The core component of the *Mandile-Schneider* algorithm is the GST, which can be constructed in linear time. The empirical measurements we have made support that *Mandile-Schneider* scales with approximately linear complexity. Future work will conduct a more thorough analysis of the algorithm's time and memory complexity, including for the worst case and the average case. We will also do additional empirical investigations of the proposed MSA algorithm on different data sets and explore alternative strategies to identify the anchors in the splitting process. Furthermore, the current implementation of trimming partially overlapping multi sub-words is not optimal and improvements could lead to further run-time gains.

For the majority of this work, we focused our attention on identifying the overlapping sub-sequences of a set of messages in order to identify the constant parts of regular expression-based message prototypes. Based on these alignments, finding suitable patterns for the non-overlapping segments in order to complete the regular expressions and evaluating their accuracy and effectiveness is a topic of an ongoing investigation.

Finally, the linear time complexity for the construction of generalized suffix trees makes them an ideal focus for further work. Particularly, we would like to investigate whether the structural information of a GST alone is sufficient to define a form of similarity between sequences that exhibits properties similar to the more commonly used edit distance.

References

[1] J. Sugerman, G. Venkitachalam, and B.-H. Lim, "Virtualizing I/O Devices on VMware Workstation's Hosted Virtual Machine Monitor," in *Proceedings of the General Track: 2002 USENIX Annual Technical Conference*. Berkeley, CA, USA: USENIX Association, 2001, pp. 1–14.

[2] C. Hine, J.-G. Schneider, J. Han, and S. Versteeg, "Scalable Emulation of Enterprise Systems," in *Proceedings of the 20th Australian Software Engineering Conference (ASWEC 2009)*, C. Fidge, Ed. Gold Coast, Australia: IEEE Computer Society Press, Apr. 2009, pp. 142–151.

[3] M. Du, J.-G. Schneider, C. Hine, J. Grundy, and S. Versteeg, "Generating Service Models by Trace Subsequence Substitution," in *Proceedings of the 9th International ACM SIGSOFT Conference on the Quality of Software Architectures (QoSA '13)*, A. Koziolek and R. Nord, Eds. Vancouver: ACM, Jun. 2013, pp. 123–132.

[4] M. Du, S. Versteeg, J.-G. Schneider, J. Han, and J. Grundy, "Interaction Traces Mining for Efficient System Responses Generation," in *Proceedings of the 2nd International Workshop on Software Mining (SoftMine '13)*, Palo Alto, CA, Nov. 2013.

[5] P. Bieganski, J. Riedl, and J. V. Carlis, "Generalized Suffix Trees for Biological Sequence Data: Applications and Implementation," in *Proceedings of the 27th Annual Hawaii International Conference on System Sciences*. Hawaii: IEEE, Jan. 1994, pp. 34–55.

[6] J. D. Thompson, D. G. Higgins, and T. J. Gibson, "CLUSTAL W: improving the sensitivity of progressive multiple sequence alignment through sequence weighting, position-specific gap penalties and weight matrix choice," *Nucleic Acids Research*, vol. 22, no. 22, pp. 4673–4680, Nov. 1994. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/7984417

[7] P. Weiner, "Linear Pattern Matching Algorithms," in *Proceedings of the 14th Annual IEEE Symposium on Switching and Automata Theory*. IEEE, Jun. 1973, pp. 1–11.

[8] E. S. Ristad and P. N. Yianilos, "Learning String-Edit Distance," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no. 5, pp. 522–532, May 1998.

[9] T. A. Sudkamp, *Languages and Machines: An Introduction to the Theory of Computer Science*, 3rd ed. Boston, MA: Addison-Wesley, 2005.

[10] J. Sermersheim, "Lightweight Directory Access Protocol (LDAP): The Protocol," RFC 4511 (Proposed Standard), June 2006. [Online]. Available: http://www.ietf.org/rfc/rfc4511.txt

[11] B. W. Watson, "A New Algorithm for the Construction of Minimal Acyclic DFAs," *Science of Computer Programming*, vol. 48, no. 2, pp. 81–97, Aug. 2003.

[12] Y. Tang, X. Lu, and B. Xiao, "Generating Simplified Regular Expression Signatures for Polymorphic Worms," in *Autonomic and Trusted Computing*, B. Xiao, L. Yang, J. Ma, C. Müller-Schloer, and Y. Hua, Eds. Springer, 2007, pp. 478–488.

[13] E. Ukkonen, "On-line Construction of Suffix Trees," *Algorithmica*, vol. 14, no. 3, pp. 249–260, Sep. 1995.

[14] M. Höhl, S. Kurtz, and E. Ohlebusch, "Efficient Multiple Genome Alignment," *Bioinformatics*, vol. 18, no. 1, pp. S312–S320, 2002.

[15] L. Chi and K. Hui, "Color Set Size Problem with Applications to String Matching," in *Combinatorial Pattern Matching*, ser. LNCS 644, A. Apostolico, M. Crochemore, Z. Galil, and U. Manber, Eds. Springer, 1992, pp. 230–243.

[16] R. Durbin, S. R. Eddy, A. Krogh, and G. Mitchison, *Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids*. Cambridge University Press, Jun. 1998.

[17] S. B. Needleman and C. D. Wunsch, "A General Method Applicable to the Search for Similarities in the Amino Acid Sequence of Two Proteins," *Journal of Molecular Biology*, vol. 48, no. 3, pp. 443–453, 1970.

[18] T. Smith and M. Waterman, "Identification of Common Molecular Subsequences," *Journal of Molecular Biology*, vol. 147, no. 1, pp. 195–197, 1981.

[19] L. Wang and T. Jiang, "On the Complexity of Multiple Sequence Alignment," *Journal of Computational Biology*, vol. 1, no. 4, pp. 337–348, 1994.

[20] N. Saitou and M. Nei, "The Neighbor-joining Method: A New Method for Reconstructing Phylogenetic Trees," *Journal of Molecular Biology and Evolution*, vol. 4, no. 4, pp. 406–425, Jul. 1987.

[21] G. Myers and W. Miller, "Chaining Multiple-Alignment Fragments in Sub-Quadratic Time," in *Proceedings of the 6th annual ACM-SIAM Symposium on Discrete Algorithms*. San Francisco, CA: ACM, Jun. 2013, pp. 38–47.

[22] A. L. Delcher, S. Kasif, R. D. Fleischmann, J. Peterson, O. White, and S. L. Salzberg, "Alignment of whole Genomes," *Nucleic Acids Research*, vol. 27, no. 11, pp. 2369–2376, Nov. 1998.

[23] D. Box, D. Ehnebuske, G. Kakivaya, A. Layman, N. Mendelsohn, H. F. Nielsen, S. Thatte, and D. Winer, "Simple Object Access Protocol (SOAP) 1.1,," W3C, W3C Note 8, May 2000, http://www.w3.org/TR/2000/NOTE-SOAP-20000508/.

[24] R. Long, M. Harrington, R. Hain, and G. Nicholls, *IMS Primer*, 1st ed. IBM International Technical Support Organisation, Jan. 2000, http://www.redbooks.ibm.com/redbooks/pdfs/sg245352.pdf.

[25] *CA Directory Administration Guide (r12.0 SP11)*, CA Technologies, 2012. [Online]. Available: https://supportcontent.ca.com/cadocs/0/g008061e.pdf

[26] J. Michelsen, "Key Capabilities of a Service Virtualization Solution," Oct. 2011, ITKO White Paper. Available at: http://www.itko.com/resources/service_virtualization_capabilities.jsp.